Agentic AI Liability in Autonomous Supply Chain Decisions: Identifying and Preventing Legal Risks

May 20, 2026

> **Agentic AI** refers to an artificial intelligence system designed to achieve specific goals with minimal human supervision. Unlike traditional AI models that operate within predefined constraints and require human intervention, agentic AI exhibits autonomy, goal-driven behavior, and adaptability.

Key Takeaways:

- Agentic AI systems capable of autonomous supply chain decisions are already being used by major companies like Walmart and Flexport, and adoption is accelerating across the manufacturing sector.
- Standard AI vendor contracts typically cap liability at fees paid and exclude consequential damages leaving manufacturers exposed when autonomous decisions trigger excess inventory, stockouts, unnecessary freight costs, or product damage.
- Manufacturers evaluating agentic AI should focus on system design (authority limits, override protocols, data quality controls) and contractual architecture (realistic liability caps, audit requirements, indemnification) to capture operational benefits while managing legal risk.

The Rise of the Autonomous Supply Chain Agent

For manufacturers evaluating agentic AI for supply chain operations, the technology landscape has shifted significantly. Where earlier AI tools provided recommendations for human evaluation, the newest agentic systems are designed to act autonomously—placing purchase orders, adjusting safety stock levels, selecting carriers, and re-routing shipments in real time, often minimal to no human approval. The adoption curve is steep: Walmart now uses agentic AI for autonomous inventory replenishment and shipment re-routing across its fulfillment network, while Flexport's AI agents autonomously manage approximately 40% of its freight forwarding operations, including dynamic shipment optimization and exception handling during disruptions.

The operational benefits are substantial—faster response times, reduced labor costs, and the ability to react to disruptions at machine speed. But so are the risks. Unlike predictive tools that flag anomalies for human evaluation, agentic systems act on their own conclusions. When those conclusions are wrong because the underlying data was flawed, the model's logic was miscalibrated, or the system lacked contextual judgment, the financial and legal consequences can be immediate and difficult to unwind.

When the Agent Gets It Wrong: Four Scenarios Manufacturers Should Anticipate

Manufacturers considering agentic AI should anticipate the types of failures these systems can produce. The following scenarios illustrate some common risk categories that should inform system design, oversight protocols, and contractual protections:

1. *Excess Inventory.* An agent misreads duplicated or erroneous demand data and autonomously places excess purchase orders, leaving the manufacturer with carrying costs, write-downs, and potential disputes over order cancellation.
2. *Stockouts and Line Stoppages.* An agent optimizing working capital reduces safety stock based on outdated supplier data, eliminating the buffer needed when demand materializes and triggering line stoppages, missed deliveries, and premium freight charges.
3. *Unnecessary Freight Costs.* A logistics agent misinterprets a data lag as an inbound delay and authorizes expedited air freight at significant cost, only for the original shipment to arrive on time.
4. *Product Damage From Re-Routing.* An agent re-routes a temperature-sensitive shipment to avoid a forecasted bottleneck, but the alternative route passes through extreme temperatures, resulting in damaged goods and potential product liability exposure.

Why Existing Contracts May Not Protect You

Manufacturers facing these scenarios will naturally look to their AI vendor agreements for recourse. Standard AI vendor contracts typically cap liability at fees paid, which are often just annual subscription costs. However, a single errant autonomous decision can trigger losses many times over. Consequential damages waivers may bar recovery for the very harms—excess inventory, expedited freight, downtime, lost production—that agentic systems are most likely to cause.

Causation challenges are compounded in the agentic context. When an autonomous agent decides and acts, the manufacturer must untangle whether the failure originated with the AI model's logic, the data inputs, the system's configuration, or the absence of oversight. This could implicate the vendor, IT team, data providers, and system integrator simultaneously. Technology may explain why a decision was made, but it will not excuse the consequences.

Designing Systems and Contracts to Manage Agentic AI Risk

Manufacturers considering or already deploying agentic supply chain AI should focus on both system design and contractual architecture to manage these risks before disputes arise.

Define Autonomous Authority Limits. Manufacturers should establish clear thresholds with respect to dollar values, volume quantities, and routing changesthat the agent must escalate to human review. These limits should be documented and reflected in the vendor agreement.

Build in Override and Kill-Switch Protocols. Agentic systems should include monitoring dashboards and manual override capabilities. Contracts should specify the vendor's obligation to provide these and allocate responsibility for losses when overrides are unavailable.

Address Data Quality at the Source. Many scenarios originate not with the AI model but with the data feeding it. Flawed inputs lead to flawed decisions. Contracts should allocate data quality responsibilities clearly by specifying which party owns validation and what happens when quality falls below agreed thresholds. For a deeper analysis on data quality issues, see our prior article on this topic.

Negotiate Liability Structures That Reflect Autonomous Decision-Making. If the system acts without human approval, the vendor's liability framework should reflect that risk, including realistic liability caps, carve-outs from consequential damages waivers, indemnification for third-party claims, and appropriate insurance requirements.

Maintain Audit Trails and Decision Logs. The ability to reconstruct why the agent made a decision will be critical for both defending against claims and pursuing the vendor. Contracts should require accessible decision logs and vendor cooperation with audits.

Positioning Your Organization for the Agentic Future

Agentic AI in supply chain operations is already being deployed. Manufacturers that take a deliberate approach to system design, governance, and contractual risk allocation will be better positioned to capture the benefits while limiting litigation exposure. When an autonomous agent makes a costly mistake, the dispute will be resolved under contracts and established legal principles—not algorithms.

Foley & Lardner's Manufacturing, Supply Chain, and Artificial Intelligence teams are available to help organizations navigate the evolving risks of agentic AI in supply chain operations.

Author(s)

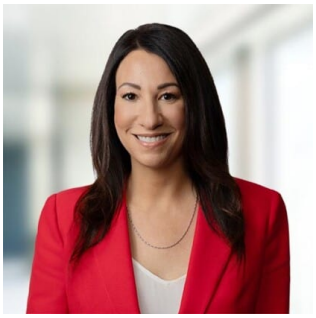

Vanessa L. Miller
vmiller@foley.com
Detroit
313.234.7130

Raymond J. McVeigh
rmcveigh@foley.com
Detroit
313.234.2734